AD_____

Award Number: W81XWH-06-2-0074

TITLE: A Clinician-Centered Evaluation of the Usability of AHLTA and Automated Clinical Practice Guidelines at TAMC

PRINCIPAL INVESTIGATOR: CAPT Emory Fry

CONTRACTING ORGANIZATION: The Geneva Foundation
Lakewood WA 98496-8687

REPORT DATE: 31 March 2011

TYPE OF REPORT: Final

PREPARED FOR: U.S. Army Medical Research and Materiel Command
Fort Detrick, Maryland 21702-5012

DISTRIBUTION STATEMENT: Approved for Public Release;
Distribution Unlimited

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED
1 March 2011	Final	29 Sep 2006–31 Mar 2011

4. TITLE AND SUBTITLE
A Clinician-Centered Evaluation of the Usability of AHLTA and Automated Clinical Practice Guidelines at TAMC

5a. CONTRACT NUMBER

5b. GRANT NUMBER
W81XWH-06-2-0074

5c. PROGRAM ELEMENT NUMBER

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

6. AUTHOR(S)
CAPT Emory Fry

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
The Geneva Foundation
Lakewood, WA 98496

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
U.S. Army Medical Research and Materiel Command
Fort Detrick, Maryland 21702-5012

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION / AVAILABILITY STATEMENT Approved for Public Release; Distribution Unlimited

13. SUPPLEMENTARY NOTES

14. ABSTRACT
~~Purpose~~: The project will demonstrate that Clinical Decision Support (CDS) material can be retrieved from a central, shared repository and executed within the MHS and civilian health information systems to improve quality of care through the use of reminders, alerts, guidelines etc. The system design and approach will decrease guideline development time and speed translation of evidence based medicine into clinical practice. It will decrease costs and enable multiple stakeholders to work in an open content/source environment to exchange clinical content, develop and test technology and explore processes in applied CDS.
~~Design:~~ Comparative study between the KMR infrastructure and capabilities developed as an open source, vendor agnostic solution for aCPG execution within AHLTA and the current DoD/MHS standard evaluating:

H1: An open source, open standard KMR and Clinical Decision Support Engine can enable organizations to share domain knowledge, expertise, and collaborate on efforts to improve patient safety and quality of care using automated clinical practice guidelines.
H2: An open standard KMR and Clinical Decision Support Engine can codify clinical practice guidelines and domain knowledge so that they can be executed without modification within a variety of runtime environments, including AHTLA and VistA.
H3: An open standard KMR and Clinical Decision Support Engine can effectively manage both generic clinical domain knowledge (guidelines, best practice, etc) and specific institutional requirements (workflow, policy, capabilities) to ensure reliable execution of a CPG across institutional boundaries.

~~Methodology~~: The KMR project architecture will leverage the emerging Federal Health Architecture (FHA) group NHIN-Connect infrastructure and its ability to provide health information exchange across a distributed network. In order to better support KMR integration and evaluation by the Office of the National Coordinator, the KMR team will coordinate and integrate the KMR project into the Agile SCRUM iterative build and test process which is in use by FHA.

15. SUBJECT TERMS
Open Source and standards based KMR and Clinical Decision Support Engine

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			USAMRMC
"U"	"U"	"U"	"UU"	41	19b. TELEPHONE NUMBER *(include area code)*

Distributed Decision Support System
&
Knowledge Management Repository

Final Report

The KMR project is supported by Award Number W81XWH-06-2-0074 administered by the U.S. Army Medical Research Acquisition Activity, 820 Chandler Street, Fort Detrick, MD 21702-5014. The opinions of the author do not necessarily state or reflect those of the Department of Defense, the Department of Veteran Affair, the Indian Health Services or of the United States Government, and shall not be used for advertising or product endorsement purposes.

Table of Contents

INTRODUCTION

The Geneva Foundation, its principal investigator, and the assembled engineering team assumed fiduciary responsibility of the "Knowledge Management Repository For Clinical Decision Support (CDS) And Proof Of Concept For Automated Clinical Practice Guideline Execution Within AHLTA" research grant in March 2009. It did so at a time when interest in clinical decision support began to receive renewed national attention and the Federal government was making large capital investments in health information technology. Clinical decision support was seen as a critical component enabling organizations to demonstrate meaningful use and achieve measurable improvements in patient outcomes. It was felt to have a central role in the challenge of patient empowerment, privacy and self-direct care, and in addressing the rising, exponential costs of healthcare delivery. Organizations were preparing to make costly investments in knowledge management infrastructure for they believed this cost would be offset, in part, by optimizing capacity and demand for precious healthcare resources.

Given this context, the KMR team received approval from TATRC/MRMC to refocus the scope of the original award from a usability study of clinical guidelines in AHLTA to investigating the creation of executable rules and workflow modules that could be shared between organizations, and that were capable of reasoning over distributed data stores from the VA, the Indian Health Service, and the Department of Defense. The team endeavored to demonstrate that this would be possible in real time using commercially available technology already made cost effective and standardized by virtue of a decade of commercial competition in the larger business market. The goal was to better define the realities of what can actually be accomplished today and articulate what still needs to be done in in order to implement a capable infrastructure for the future.

SUMMARY OF DELIVERABLES

Deliverable #1: Project Documentation and Engineering Plans
Deliverable #2: Functional Requirements, Use Cases, and Academic Review
Deliverable #3: System Design Specification and Academic Review
Deliverable #4: User Documentation
Deliverable #5: System Test, Integration, and QA Plans
Deliverable #6: Metadata, Terminology/Ontology, & Content Specifications
Deliverable #7: KMR Repository and Content Management System
Deliverable #8: Rule and Guideline Runtime Engine
Deliverable #9: System Administration Portal and Performance Monitoring Tool
Deliverable #10: Point of Care CDS Tool for AHLTA
Deliverable #11: Guideline / Rules Workbench and Authoring Tools
Deliverable #12: Content and Executable Clinical Guidelines
Deliverable #13: Academic presentation & demonstration of runtime deliverables
Deliverable #14: Academic review of current CDS 'state of the art'
Deliverable #15: Academic panel discussion on future state CDS requirements
Deliverable #16: Academic review of completed KMR Project
Deliverable #17: Academic outcome & usability evaluations

BODY

Given the new Statement of Work, the KMR team was faced with the challenge of defining the general scope and depth that was appropriate for the project. A narrowly defined scope with a deep and full implementation highlighting technical refinement and polish for select use cases would have been possible, but might not have illustrated how important a comprehensive and methodically executed architectural strategy was to realizing the full potential of clinical decision support. Therefore the team deliberately chose a broader, prototypical approach emphasizing service and component orchestration vice the particulars of a single perspective. While certain KMR concepts were later fully developed and implemented in a production setting by the Military Health System (MHS) using additional resources from TATRC, the majority of our deliverables remain at the research and prototype level of development.

1. Deliverable: Project Documentation & Engineering Plans

The KMR engineering methodology was a combination of "Waterfall" and "SCRUM" development perspectives. This approach to product development combines agile best practices within a simplified, more traditional Waterfall strategy. A project starts with an abbreviated Waterfall scoping and design phase to establish prioritized requirements, build consensus amongst stakeholders and define high-level system design objectives.

This more traditional phase is then followed by rapid, iterative development using the SCRUM methodology. Through this collaborative process, products are iteratively developed with close collaboration with stakeholders, providing transparency into development progress, earlier delivery of functionality, and faster realization of value. It is well suited for performing proof of concept evaluations for new project ideas and technologies while still delivering scalable, enterprise capable architectures for projects selected for deployment. We conducted 27 such iterative cycles or "sprints".

SCRUM is fundamentally a simple methodology, typically executed by a small team consisting of 2-4 software engineers, a quality assurance specialist, and a SCRUM manager. It does NOT, however, conform to traditional waterfall metrics and Gantt charting reporting – the process uses concepts such as "sprint velocity", "burn-down charts", and "story points". Consequently, SCRUM teams usually use specialized agile software tools designed to accommodate the highly flexible and variable scheduling and tasking that the SCRUM process encourages. Our team used an on-line tool called Target Process that we use successfully in all our development efforts.

Representative user stories, charts, and engineering tasks that were used to track engineering progress online have been exported and can be found in the "Engineering Plan" folder on the Subject Data DVD.

2. Deliverable: Clinical Decision Support Requirements

A primary focus of our research was establishing a comprehensive set of requirements for clinical decision support in general and KMR in particular. During the development of these scenarios, several major themes emerged, most notably the need to define the

semantic and structural underpinnings for rule authoring and execution. At a base level, sharing of rules implies a repository of some kind that can be searched efficiently to discover artifacts of interest using terminology that is unambiguous and granular. The rules themselves also needed to be standardized in terms of structure and semantics. Current inference technology demands that the rules they execute be tightly and irrevocably bound to the particular structure and semantic meaning of the data they are asked to analyze. If that same rule is shared with an organization whose data differs from the structure used during the authoring phase, the rule will fail to execute. More importantly, if an organization's data structure is identical, but its semantics differ, the shared rule may indeed execute, but with unpredictable, and potentially dangerous, consequences. Finally, in order to share and execute clinical rules safely within different organizations, additional metadata regarding the authors original intent, and the clinical context that the rule is assumed to operate in must accompany the rule itself, least a structurally and syntactically correct rule results in a decision that is inappropriate in a new environment. For example, a rule for an adult head trauma patient recommending an MRI would be inappropriate in a pediatric setting with only a CT scanner. These requirements for semantic integrity of data across organizations proved to be a major determining principle to the KMR design, and are discussed in more detail in Section 6.

A second category of requirements relates to the specificity of the rule itself. Computable guidelines represent, by their nature, best clinical practice for general populations. Providers and patients were quite vocal in our focus groups regarding their frustration with current CPG's that apply generalities to a particular individual and are prone to alerting providers of recommendations that are clearly inappropriate given a particular patient's past medical history. An effective clinical decision support system must define or infer inclusion and exclusion criteria before it is indiscriminately applied. It should have the capability of accommodating its recommendations to the patient's individual directives. Similarly, rules and guidelines are usually created and approved at an organizational level and subsequently applied to groups of providers. These recipients express great frustration when they see their care as being dictated by the clinical decision support system or perhaps even dismissive of the skill and experience they bring to patient care. While they understand the benefits of evidence-based guidelines in reducing unnecessary variance in clinical practice, they justifiably resist any process that they perceive as reducing their ability to practice the "art of medicine". The consensus opinion was that CDS systems should provide individual providers with tools to supplement broadly applicable institutional rules with personal, handcrafted provider specific knowledge. The technical approach to striking a balance between reducing variance at the organizational level and individual autonomy and expressiveness was the principle driver for the design of the core runtime rule engine and its companion CDS workbench.

Finally, our functional scenarios repeatedly highlighted the need for near-real time performance, particularly for inpatient clinical scenarios. This performance was expected not only for simple data evaluations, but also for rules utilizing complex statistical and/or historical trend analyses. Given the large volumes of data that such decisions might require, the quality of service requirements of our current EMR implementations, and the

operational constraints imposed by the need to reason over distributed data stores, these functional requirements in particular lead to an architectural design for the CDS system unlike any prior CDS implementation. This unique approach is described in more detail in subsequent sections.

The requirements process elaborated many more concepts and requirements than were in scope for the project – this over reach was both deliberate and required. Only by understanding the broad breadth of possible CDS uses could the team adequately anticipate the larger architectural requirements that loom on the HIT horizon. The time and energy devoted to the requirement phase enabled the design of comprehensive, standards-based SOA architecture that is adaptable to any healthcare organization willing to expose its legacy data, and can satisfy the core requirements for truly distributed decision support.

Deliverable is located in the "Requirements" folder on the Subject Data DVD.

3. Deliverable: System Design Specification & Academic Review

Our functional scenarios and requirements process proved invaluable in guiding us towards an architecture that was semantically constrained, adaptable, and able to "normalize" the structure of any data it reasons with. It needed to be able to collate patient information from across multiple organizations and it had to provide a level of performance never before expected. This architecture had to be configurable to both patient and provider preferences, and had to provide the capability to change behavior.

We believe the final KMR design delivers all these capabilities, and does so in large part because it was able to leverage and supplement the components deployed in the basic FHA-CONNECT architecture for the Nationwide Health Information Network (NHIN). CONNECT implements a flexible, open-source gateway solution that enables healthcare entities to connect their existing health information systems to the NHIN. It is uses service-oriented-architecture design principles and web service interfaces. This architecture enables individual components to be replaced by custom solutions as long as they adhere to the defined web service interface specifications. It also allows implementations to be hosted on different hardware and software platforms, as well as services to be implemented using different programming languages.

KMR is designed as supplemental components deployed on the basic CONNECT architecture, although in the KMR implementation CONNECT is configured somewhat differently than the default release. The overall architecture of the system can be logically broken down into two sections: the CONNECT gateway with corresponding services and components, and an adapter with corresponding adapter services and components. The gateway connects the existing health information system(s) of an organization to the NHIN. Gateway services provide mechanisms for receiving messages from the NHIN and passing them to the adapter Clinical Decision Support (CDS) service, as well as for receiving CDS messages from the adapter and sending them to the NHIN. In addition to supporting the core NHIN services, components in the gateway also provide services to manage NHIN connection endpoint URL data, patient correlation, and a variety of other

supporting services. The CONNECT adapter comprises the software that interfaces the existing health information system of an organization to the CONNECT gateway.

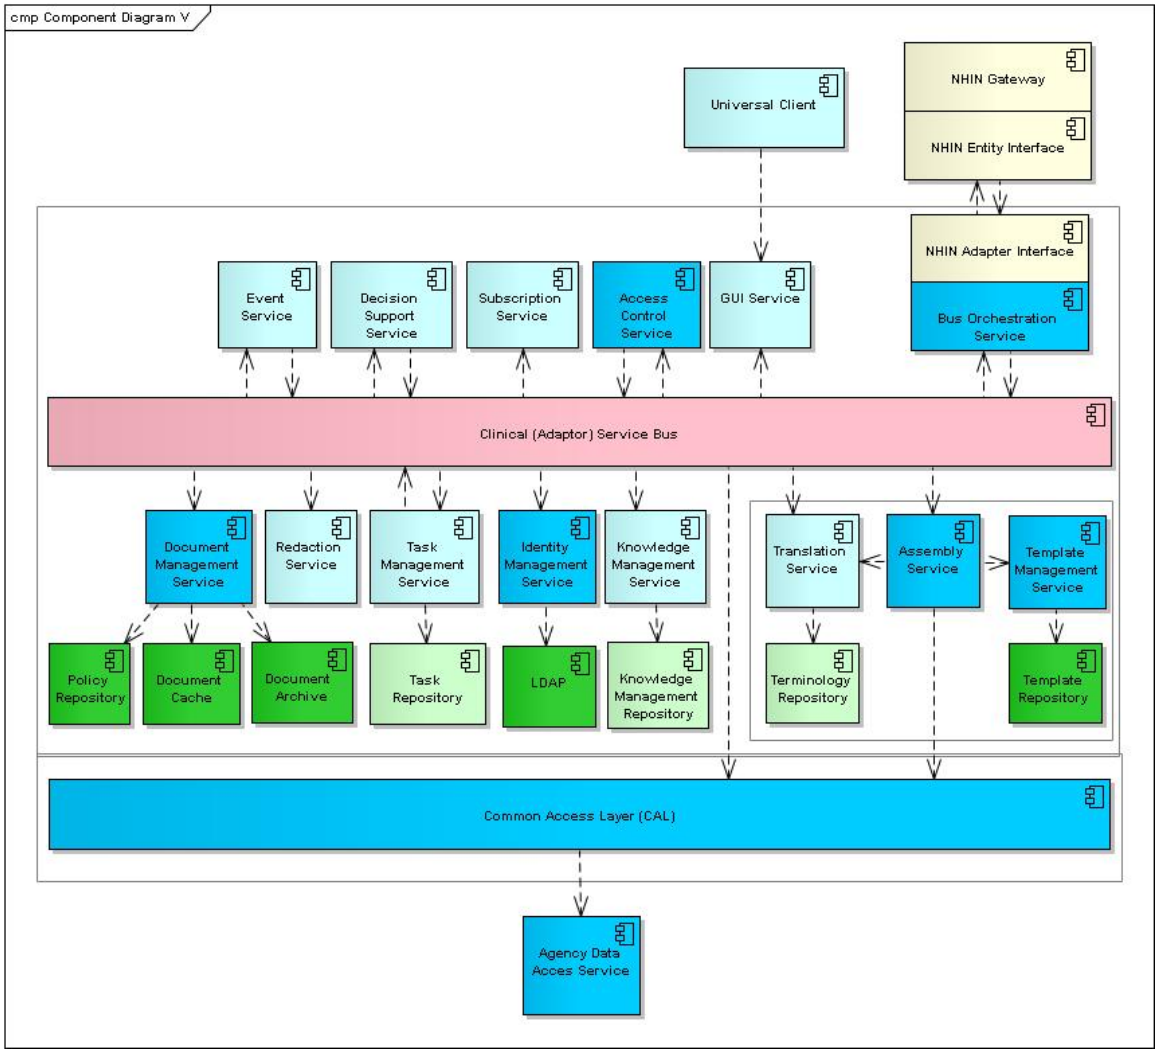

Figure 1: Component Overview

Service #1: Subject Discovery

In order to share patient data between connected organizations, it is necessary to have mechanisms to match patient identities in the absence of a single national identifier. Subject Discovery represents a set of services that meets this need by providing the mechanism for locating patients, or "subjects", based on demographic information. These services provide the ability for one organization to determine whether other organizations have records for a given patient by submitting a set of demographic identifiers that organizations can use to match against their own master patient indices. In the KMR implementation, the Subject Discovery Service remains unchanged.

Service #2: Query for Documents
The Query for Documents service provides the mechanism by which one organization locates electronic health information on the NHIN associated with a specific subject. The service allows one organization to acquire a list of documents for a given patient that may exist at another organization, based on a set of search criteria. The service can be viewed as receiving query requests from the NHIN to which it must respond, and sending queries to organizations on the NHIN in order to locate a patient's health information. In the KMR implementation, the Query Service was modified to be a fully automated broadcast that can be executed by a rule, instead of requiring each site to be manually selected and queried.

Service #3: Retrieve Documents
The Retrieve Documents service provides a mechanism to retrieve the electronic health information on the NHIN. It is used in conjunction with the Query for Documents service, which returns a list of document references that Retrieve Documents uses to retrieve patient records. The service can be viewed as receiving document requests from the NHIN to which it must respond, and sending requests to organizations on the NHIN in order to retrieve patient health information. In the KMR implementation, the Retrieve Documents service was modified to be fully automated and executed by the rule engine.

Service #4: Query Audit Log
The Query Audit Log service provides the mechanism by which audit data associated with accessing health information on the NHIN is exchanged so that consumers and privacy or security officers can account for who has had access to what information for what purpose. The service allows one organization to request an audit log, meeting certain search criteria, from another organization. The service can be viewed as receiving query requests from the NHIN to which it must respond, and sending queries to organizations on the NHIN in order to receive and potentially view audit log entries. In the KMR implementation, the Query Audit Log service remains unchanged.

Service #5: Authorized Case Follow-up
Pseudonymization is the process of removing the association between a data set (e.g., protected health information or PHI) and the subject of that data (e.g., a patient) by removing identifying information, and adding an association between the data set and one or more alternative identifiers, or pseudonyms. Re-identification is the process of obtaining the association between a pseudonym and the original subject of that data set (e.g., re-associating PHI with the patient). Pseudonymization may be required for a number of reasons. For example, there may be a need to report health information to a public health agency for surveillance purposes in which the identity of the subject is not needed. Likewise, re-identification may be required if an authorized individual, such as a public health official, must investigate a potential public health issue with proper legal authorization by gaining more information from the longitudinal health record of the individual known only by their pseudonym. Authorized Case Follow-up represents the services for re-identification. It does not include the services, algorithms, or specifications for pseudonymization. In the KMR implementation, the Authorized Case Follow-up service remains unchanged.

Service #6: Health Information Event Messaging
CONNECT serves two primary workflows: an organization-initiated subscribe message / NHIN notify message, and an NHIN-initiated subscribe message / organization notify message.

In the first workflow, an existing system serviced by CONNECT initiates a subscription to the entity interface. The gateway records this parent subscription in the subscription repository, and contacts other appropriate NHIN-enabled organizations with the request to receive notifications of available data. The remote organization's response contains a child subscription reference, which is recorded in the subscription repository associated with the parent subscription request. When a notify message is received from a remote NHIN-enabled organization, the notify message is matched to its' child subscription. The child's parent subscription is retrieved, and used to build a new notify message which is passed to the adapter for processing.

In the second workflow, a remote system on the NHIN initiates a subscription to CONNECT, which records this subscription in the subscription repository. Based on configuration, CONNECT will then send a child subscription on to the existing system serviced by CONNECT. When a notify message is received from the system via the CONNECT adapter, the notify message is matched to the subscription, either by subscription reference if provided, or by criteria matching. Once a match is found, the notify message is sent to the appropriate remote NHIN organization(s).

The HIEM service includes services for both workflows, to manage subscriptions and process notifications, both to and from the NHIN. The level of subscription information provided to the existing system that is serviced by CONNECT - i.e., whether no notification is made, whether notifications are copied to the system, or whether child subscriptions are created - is configurable. In the KMR implementation, the HIEM service remains unchanged.

Service #7: Master Patient Index
The Master Patient Index is based on the open source Mural project. Mural includes the following set of core components. In the KMR implementation, the Master Patient Index service remains unchanged.

Service #8: Document Registry and Repository
The XDS.b document registry and repository are based on an open-source implementation hosted by the National Institute of Standards and Technology (NIST) in support of Cross-Enterprise Document Sharing (XDS) and related profiles published by Integrating the Healthcare Environment (IHE). The IHE XDS profile describes how to exchange clinical documents for patient care. In the KMR implementation, the Registry and Repository service was extended significantly to better comply with the XDS standard, and to accommodate the requirements for the Alert Repository.

Service #9: Access Control Service
Access Control Service is a new service provided by KMR to provide fine grained access control to system objects. The Policy Engine takes responsibility of determining whether a message should be processed by CONNECT, regardless of the direction the message is being moved - whether it is inbound from the NHIN or outbound to the NHIN - and thereby enables an organization to apply policies to all messages. Policies may be patient-specific (e.g., based on the patient's consent for a specific information exchange) or organization specific (e.g., based on hours of operation, user role, etc.).

Service #10: Audit Log
For our base release of CONNECT, the Audit Log is based on a simple implementation developed as part of the CONNECT development effort. There are no plans at this time to replace this Enterprise Service Component with any other open-source implementation. In the KMR implementation, the Audit Log service remains unchanged.

Service #11: organization Service Registry
The Service Registry interface specification provides for registry servers that enable NHIN-enabled health organizations to discover the existence and connection information for other NHIN-enabled health organizations, utilizing UDDI. In the KMR implementation, the organization Service Registry remains unchanged.

Service #13: Document Assembly Service
Whenever differing systems are integrated, there is often a need for services that transform data types from one system into data types needed by the other. Such transformations are required regardless of whether the exchange is document or message based. The Document Assembler Services, provided by KMR, are new services and components necessary to produce HITSP CDA compliant documents and manage the metadata necessary for transmission through the NHIN CONNECT Gateway if required. The Document Assembler Service ensures that “facts” from one organization can be exchanged with another so that a rule engine can reason over a distributed collection of clinical data.

Service #13: Template Repository
The Template Repository is a new component provided by KMR to manage requests and responses for schemas, transforms, and metadata regarding required data access calls for the Document Assembler to build standards based artifacts for healthcare information exchange.

Service #14: Decision Support Service
The Decision Support Service is a new service provided by KMR to execute analytic operations on behalf of other clinical applications and systems. The results of a rule evaluation are then passed back to the invoking agent and/or forwarded to additional recipients as required.

Service #15: Knowledge Management Repository & Service
The KMR Repository stores and indexes one or more decision support knowledge modules in it's' repository. It exposes required metadata and descriptors (e.g., text descriptions, clinical contexts, vocabulary requirements, etc.) to support distributed service discovery, invocation, and sharing of clinical rules. It is a new database and service provided by KMR.

Service #17: Presentation Service
GUI Services are new KMR services providing a single point of entry for all GUI development. By providing a catalog of published API calls and return values, Graphical User Interface developers are freed from having to understand the vast amount of Web Services, databases and other technical details of the KMR System. This GUI Services layer also allows us to aggregate multiple system calls into a API call where applicable. For example a simple HTTPS GET call can check every interfaced sub-system in the application returning a known set of XML or JSON formatted data to be used as the front-end developer sees fit. Developing a GUI Services layer also allows us to enforce system security while implementing role based access control from a single point.

Service #18: Common Access Layer
The Common Access Layer service is a new proxy service provided by KMR to shield adaptor services from the particulars of the enterprise data model. It provides an internal interface for invoking implementation specific data calls and maps the returned results into standard data objects. These data objects are defined by the constraints to the CDA Schema described in the C83 Content Module Specification. The end result is that data access calls produce standard based data objects upon which other adaptor services can rely upon to be structurally and semantically consistent regardless of the particulars of the implementation.

Service #19: Event Service
The Integration service is a new service provided by KMR delivering a variety of connectors and listeners to consume relevant healthcare messages, alerts, and data triggers. The service ensures that clinical events are detected, processed, and forwarded for CDS evaluation.

Service #20: Task Service
Clinical Decision Support Knowledge Modules execute logical operations and generate notification messages, alerts, tasks and trigger subsequent workflows as output. The Notification Service is a new service provided by KMR to take this output and ensure the intended recipient is notified. The service will also track notification acknowledgements and handle escalation of notifications as specified by the rule author. A Notification Repository will be used to determine who gets the result and with what protocol.

Service #21: Redaction Service
The Redaction Service is a new service provided by KMR to remove information from a patient's medical document. Currently, patient's can opt-in or opt-out of a tool that allows providers to gather the patient's medical information. The idea with the Redaction Service

is to allow the patient to designate specific sections to be excluded from their medical information document. The Redaction Service will be designed to accept C32 medical documents. Along with the document will be a list of section identifiers that need to be redacted from the document. The result of the service will be a valid C32 document with the appropriate sections removed. In place of the removed sections will be a notice that says, “This data was masked per patient consent directive".

In general, the changes and additions to existing CONNECT services ensure that orchestration and workflow management of the entire middle tier resides squarely within the adaptor. From this system perspective, the gateway becomes a specialized "client" for ensuring NHIN compliant system behavior when and where it is required to provide the rule engine with access to structured data anywhere across the network. KMR conceives of the adaptor as a basic SOA bus for healthcare, an infrastructure for delivering the data and services required for advanced clinical decision support and workflow optimization.

The challenge of integrating reasoning and process orchestration into this scalable, national reference architecture led to an in-depth discussion regarding the strengths and weaknesses of several approaches. A rule evaluation is inherently a stateless task in which the inference engine evaluates only the facts presented to it and retains no memory of its decision upon completion. Workflow, however, is by nature stateful and the engine is required to maintain some notion of time and process state. It is rare to see a single engine incorporate both workflow and rule functionality; these capabilities are most often discrete and separate given the technical difficulties inherent in building such engines.

Separating workflow from inference, however, introduces significant overhead when attempting to orchestrate the complex interplay between process flow and rule logic that is typical of clinical guidelines. Historically, CDS researchers have either created custom, nonstandard, engines to investigate these more ambitious requirements or have been restricted to focusing either on process flow or rule evaluations. The result is that the medical domain expert, who unconsciously and freely alternates between process flow and logical inference, often perceives current systems as awkward or incomplete. Our domain experts repeatedly articulated the need to develop a system that made the implementation of these cognitive models more intuitive and transparent. The discordance a user experiences when forced to use systems designed to emphasize one or the other paradigm were investigated and addressed in the design of the clinical decision support service described more completely below.

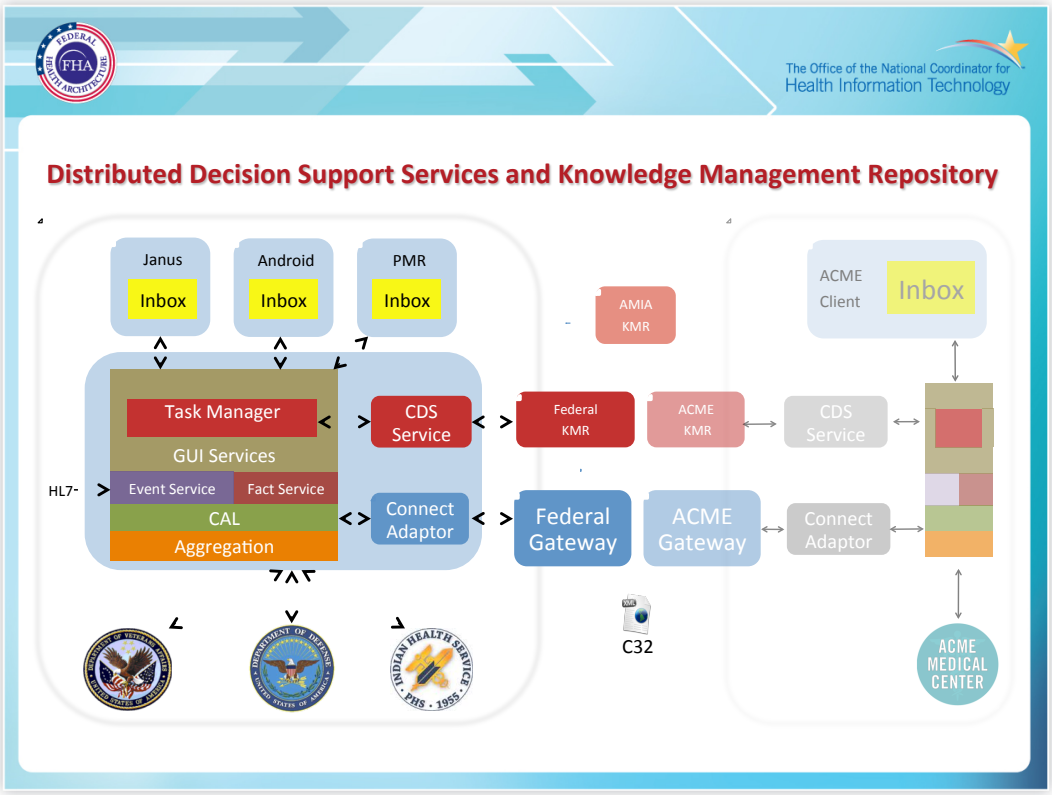

Figure 2: Overview of Cross-Domain Interactions

Deliverable is located in the "Design" folder on the Subject Data DVD.

4. Deliverable: User Documentation

Build and installation documentation is located in the "User Documentation" folder on the Subject Data DVD. Given the engineering focus of this project and the prototypic nature of the web applications, only limited, non-technical end-user guides are provided.

5. Deliverable: System Test, Integration, & QA Plans

System test, integration and functional validation were tracked online during the SCRUM process as QA user stories and tasks in Target Process. Representative user stories, charts, and engineering tasks that were used to track QA progress have been exported and can be found in the "Engineering Plan" folder on the Subject Data DVD.

6. Deliverable: Metadata, Terminology & Content Specifications

A central theme behind the entire KMR approach and architecture is that it is based on standard medical vocabularies and taxonomies. The Common Access Layer (CAL) is the foundation of this effort, the core component that implements the majority of the platform's metadata and vocabulary framework. The CAL specification enables all subsequent capabilities and ensures that KMR clinical decision support services can be

layered on any medical information system willing to expose its data. CAL is set of reference interfaces, based on HL7 version 3 standards and vocabularies that normalize the structure and semantics of legacy data into a standards-based canonical model, shielding the middle tier from the particulars of the enterprise data model. CAL provides middle tier services with an internal interface for requesting data and maps the returned results into standard data objects. These data objects are defined by the HL7 messaging standard and the constraints to the CDA Schema described in the C83 Content Module Specification. The end result is that calls for historic or legacy data produce standard based data objects upon which other adaptor services can rely upon to be structurally and semantically consistent…regardless of the particulars of the implementation. The following is a high-level diagram of the Common Access Layer Service.

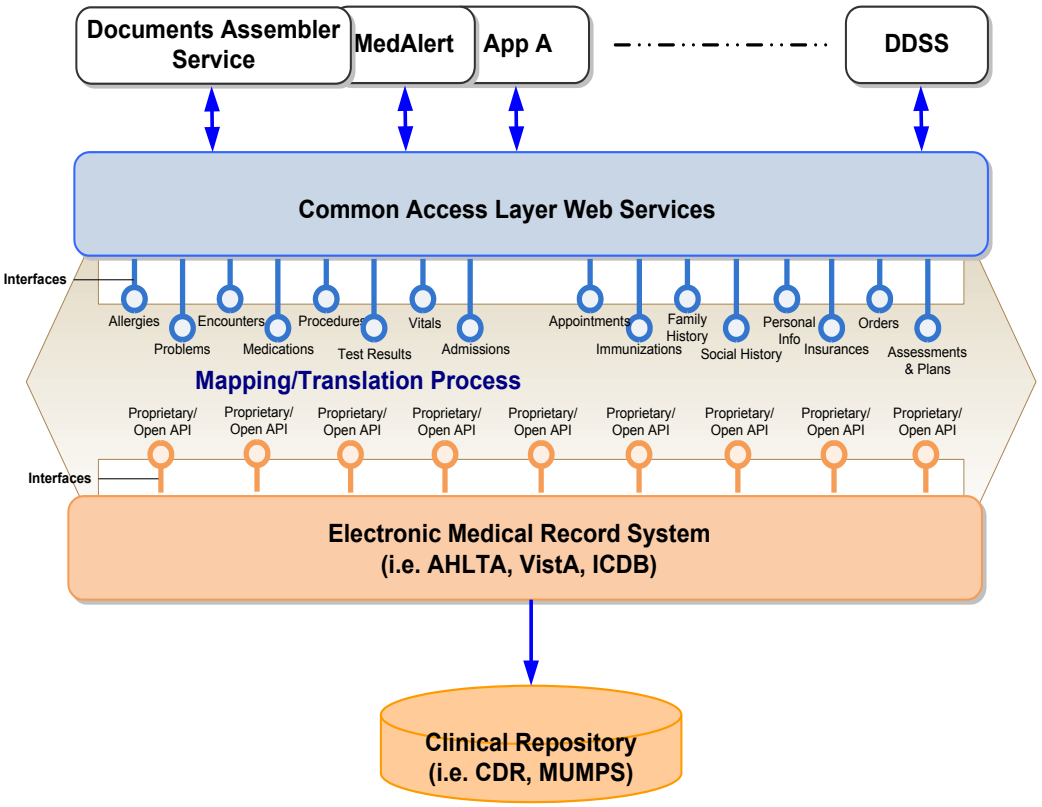

Figure 3: High-level view of Common Access Layer Service Architecture

.Event Service

The Event Service provides a variety of connectors and listeners that enable KMR to consume relevant healthcare messages, alerts, and data triggers. It is similar in concept to the CAL service, only it is responsible of intercepting real-time event messages generated, for example, by lab / radiology equipment or other transactional systems that maybe in use within the organization. It is currently configured for HL7 messages and C32 NHIN documents, but a variety of different messages types can be handled. After receiving the message, Event Service parses the payload and delivers "fact" objects to the rule engine that have the identical semantic and structure as those derived from the legacy

electronic medical record by CAL. If a C32 is received, its data will similarly be consumed and transformed.

Fact Extraction Framework

In order to ensure enterprise performance and scalability, commercially available workflow and rule engines are carefully tuned to specific data and structural characteristics. When these engines are used with clinical data or architectures that are not consistent with these requirements, their performance suffers and scalability is adversely affected. For example, most production business rule engines utilize the Rete algorithm to ensure adequate performance in environments that require the engine to reason over millions of facts using knowledge bases consisting of potentially hundreds of thousands of rules. When confronted with a typical clinical "fact" that has been modeled using HL7 RIM 3.0, these engines struggle with the highly nested and recursive nature of the standard. The Rete algorithm they employ is severely handicapped and they lose any performance or scalability advantage they might have over nonstandard proprietary research solutions.

It became apparent that the "facts" used to deliver clinical decision support had to be optimized and structured appropriately if commercial engines were to be utilized. The lack of a ballot approved HL7 object model, designed for run-time implementations, was seen as major omission, and remains a critical impediment to the national CDS agenda. The Fact Service was developed to extract the critical clinical data from these canonical objects and optimize them for runtime. This allow the system to retain the massive scalability that commercial rule engines are capable of while still maintaining the highest level of conformance to the HL7 standard when extracting data from the legacy system. The KMR team devoted significant effort in the design of the Fact Extraction Framework, and participated actively in OASIS and the HL7 Virtual Medical Working Group to communicate these implementation specific impediments.

Task Manger

The final area where KMR devoted considerable time and engineering effort was in the metadata and semantics needed for rules to adequately articulate the process and orchestration requirements that are implied by virtually all comprehensive clinical care plans. These requirements are rarely made explicit given the traditional focus on articulating rule logic. For example, it is not currently possible to accurately encode the difference between delegating and transferring responsibility for a task within a healthcare setting. When delegating a task, responsibility for the task remains with the original actor. When responsibility is transferred, the subsequent consequences of that task no longer apply to the original actor. While this may seem rather esoteric technically, it is actually a qualitative distinction critical to many common scenarios such as attending/nurse/student interactions or transfer of care situations.

Task Manager is the service provided by KMR to take the results of a rule evaluation and ensure that the intended recipient is notified. The service is essentially a service endpoint look up that accepts "task objects" from the rule engine and delivers the metadata required by the invoked service. Task Manger has the ability to write orders, book

appointments, send email/alerts/SMS messages, and register patients in our reference Disease Registry system. Each of these abilities is firmly rooted in a standards-based implementation using industry standard APIs.

The combination of the CAL Service, the Event Service, and the Fact Service ensures that the CDS rule engine has a semantically consistent fact collection over which to reason, regardless of whether the data comes from a legacy database, from the NHIN, or are delivered as an HL7 messages in real time. These three services ensure the metadata, terminologies, and content of the clinical data being analyzed is standards-based and optimized for run-time.

Deliverable is located in the "Project Source Code" folder on the Subject Data DVD.

7. Deliverable: KMR Repository & Content Management System

The KMR Repository stores and indexes one or more decision-support knowledge modules in its database. It exposes the metadata and descriptors (text descriptions, clinical context, vocabularies, etc.) required to support discovery, invocation, and sharing of artifacts between organizations. The repository maintains all the reference vocabularies required to fully annotate rule objectives, requirements, facts utilized, etc. so the rules can easily be discovered and retrieved by the functional community. It also provides two other important services. It supports, the rule development and governance process by maintaining the lifecycle management metadata required during the evolution of a clinical concept to a fully vetted and production ready clinical practice guideline. Its role-based access control mechanisms ensure that organizations have the flexibility to articulate whatever governance process is deemed appropriate.

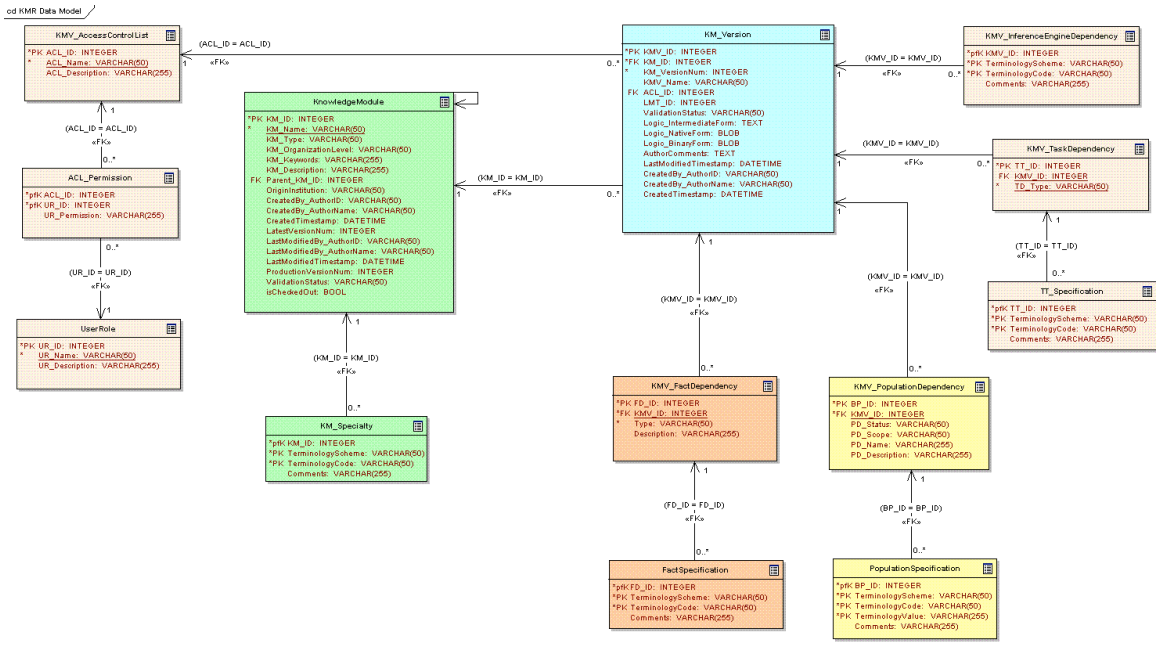

Figure 4: KMR Repository Schema (Illustrated Without Vocabulary Reference Tables)

A deceptively simple entity-attribute schema maintains fully versioned rule artifacts against which a limitless number of semantic cross-references can be maintained. These vocabulary reference tables are used during the rule-authoring phase to semantically constrain the fact types exposed to rule authors.

KMR rules, workflows and other artifacts are stored in the Knowledge Management Repository and are searchable using the Repository Search Application. This tool exposes all the meta-data used to annotate and cross-reference the content so that rules can be easily located and retrieved. Users may search by specialty, patient acuity, rule objective, disease, facts evaluated, age range, etc. Each of these search parameters is tied to the particular vocabulary that the CDS Workbench automatically enforces during the authoring phase, for example, LOINC, CPT, etc. As a web tool, it enables an organization to make its work product available to others using flexible, fine grained, role-based access control. Each site can choose precisely what artifacts are exposed and what user credentials must be presented to gain access.

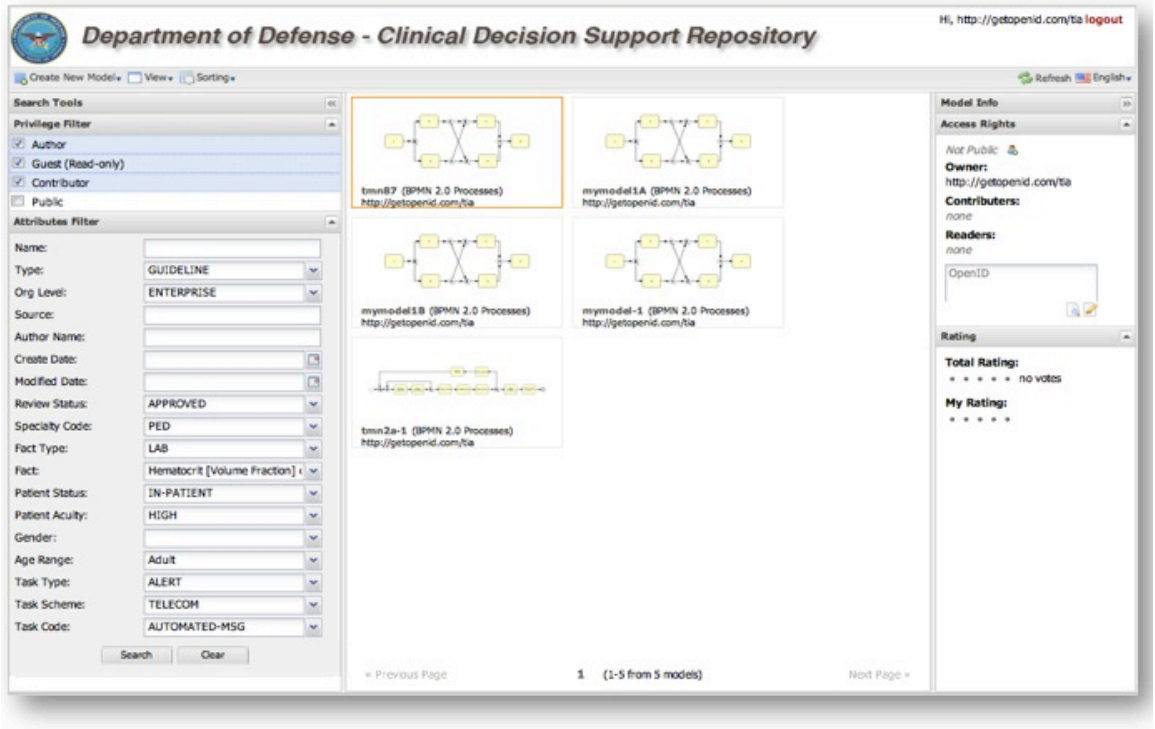

Figure 5: Rule Repository Search & Retrieval Tool

Knowledge Management Services consists of the components used to store rules and other computational artifacts in the KMR Repository, and to dynamically load and unload rules into a patient's session during run-time. It also ensures that all rule meta-data is appropriately managed and synchronized with the reference vocabularies and taxonomies chosen by the organization.

Deliverable is located in the “Project Source Code” folder on the Subject Data DVD.

8. Deliverable: Rule & Guideline Runtime Engine

The Decision Support Service (DSS) is the actual inference engine provided by KMR to execute analytic operations on behalf of other clinical applications and systems. The results of a rule evaluation are then passed back to the invoking agent and/or forwarded to additional recipients as required. DSS, its Knowledge Management Repository (KMR) and its companion services provide an infrastructure for the real-time evaluation of clinical data and the notification of the appropriate people or systems of the results.

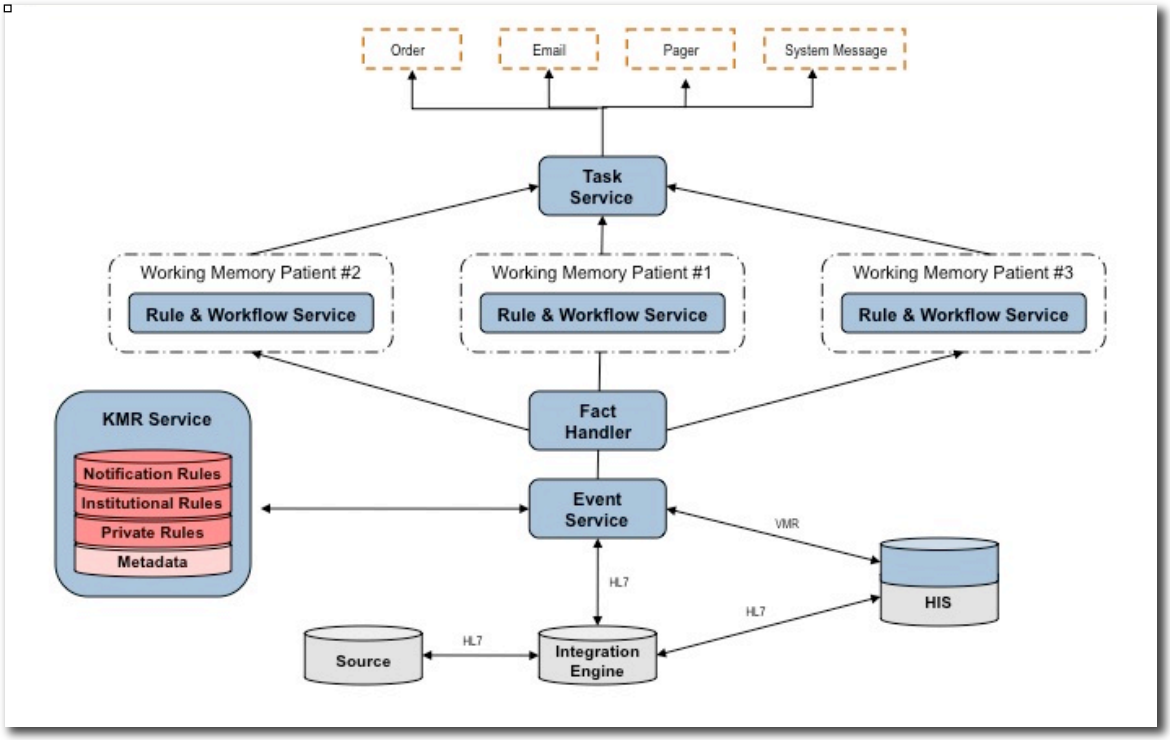

Figure 6: Individualized Rule Sessions – One Per Patient

The team chose the JBoss Drools engine to implement this CDS infrastructure. This tool is unique among commercial offerings in that it can execute process flow and rule evaluations within the same instance of the engine… using the same syntax. This capability has tremendous advantages with respect to service management, nearly seamless integration between process and inference cognitive models, and fewer resources/skill sets being required to support a production deployment.

The Drools engine is open-sourced under a very business friendly copyright license and has an active community of developers and implementers. Access to its source code and a vibrant, global network of collaborators, proved invaluable in implementing many core KMR requirements. For example, the requirement for rules highly customized to particular patients and to specific providers led to a design in which every patient is given their own instance of the rule engine in which individualized knowledge modules can be deployed. This approach is unique in the literature, but is extremely resource intensive. For a facility with a finite number of beds, controlling even several hundred inpatient

sessions is not of particular concern. On the outpatient side, the resource management problem is exponentially more complicated. The solution was to extend the engine's native session management capabilities with what has become known as Drools Grid, a grid enabled version of Drools that enables a system to dynamically instantiate sessions remotely, serialize to disk on demand, and upon de-serialization, load balance across an array of machines. Such development would simply not have been possible with a closed source product, or at least prohibitively expensive.

In addition to individualized sessions, the KMR design delivers a second major innovation… the concept of a stateful, in memory, persistent store of all historic facts that might be needed when evaluating new data. When a patient accesses care for the first time, the CDS service retrieves a large collection of historic data from the EMR and stores it in the rule session as a virtual medical record or vMR. While this initial load and set up incurs a transactional penalty on the database, virtually all-subsequent evaluations avoid additional disk access because the data is already resident in memory. Any new clinical data that comes across the wire is readily consumed and can be immediately evaluated in the context of the patient's past medical history. This design enables rules requiring statistical analyses, trend evaluations, or other complex operations to occur almost instantaneously. As demonstrated at HIMSS 2011, the performance of the KMR system is outstanding and is believed to be scalable, given adequate memory resources, to millions of patient facts and knowledge bases with tens of thousands of rules.

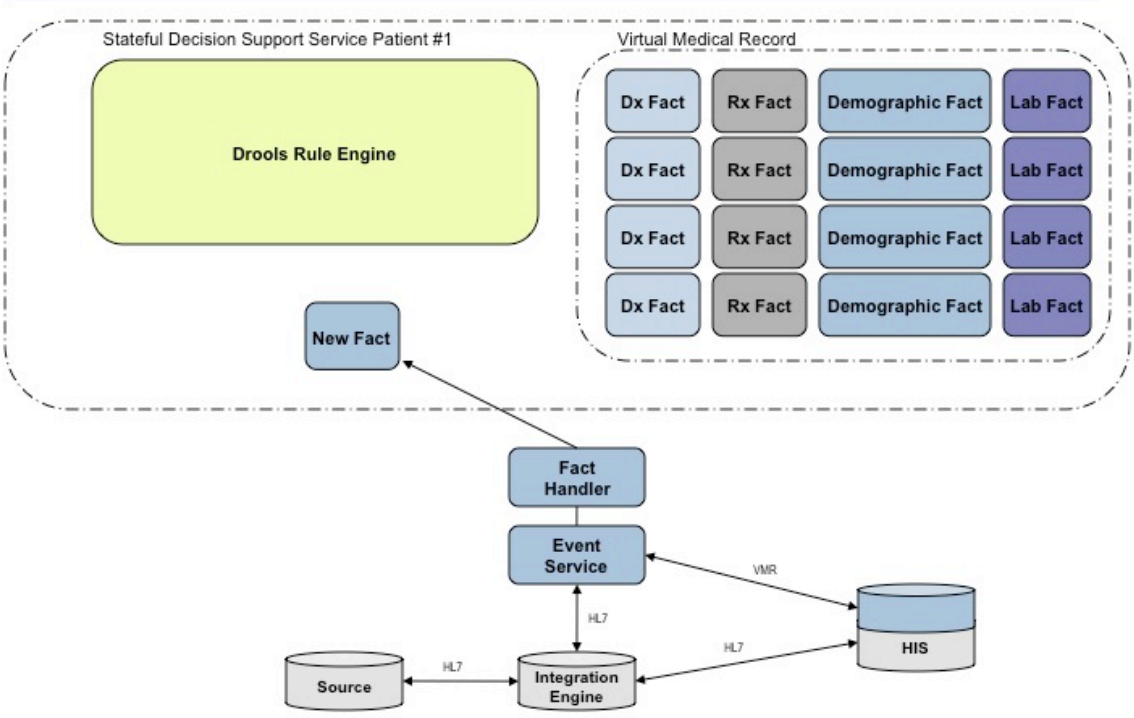

Figure 7: Virtual Medical Record

Deliverable is located in the "Project Source Code" folder on the Subject Data DVD.

9. Deliverable: Point of Care CDS Tool for AHLTA

The KMR team and its partners implemented several applications identified during the requirements process as being particularly desirable.

Universal Inbox

The Universal Inbox (previously referred to ad MedAlert) is perhaps the most important of these in that serves to unify multiple disparate workflows within a single, centralized tool. Using a familiar Microsoft Outlook metaphor, Universal Inbox receives messages of all types, collates them within a single queue, and then utilizes an extensible array of plug-ins to display each message type in a way that is visually appropriate for the information being conveyed. The team implemented several plug-ins including an NHIN C32 viewer and a clinical CDS alert parser able to dynamically render rule-defined action buttons that allow the user to execute the particular task recommended by the alert. Many other possibilities exist, include survey tools, file utilities by which patients can manually upload scanned documents, or graphing components used to visualize medical device data obtained, for example, from a glucometer. This plug-in architecture ensures that the Universal Inbox can be enhanced with new capabilities as the number and type of messages that might be aggregated within the mail queue grows over time.

The Universal Inbox itself is a module designed to be integrated within a wide variety of applications. Using only a limited subset of metadata, specifically the unique provider and patient ID, the Inbox is able to establish appropriate context and render a role-based collection of messages. The Universal Inbox has been successfully incorporated into the AHLTA client, the Indian Health Service RPMS client, and the VA Janus Provider Portal.

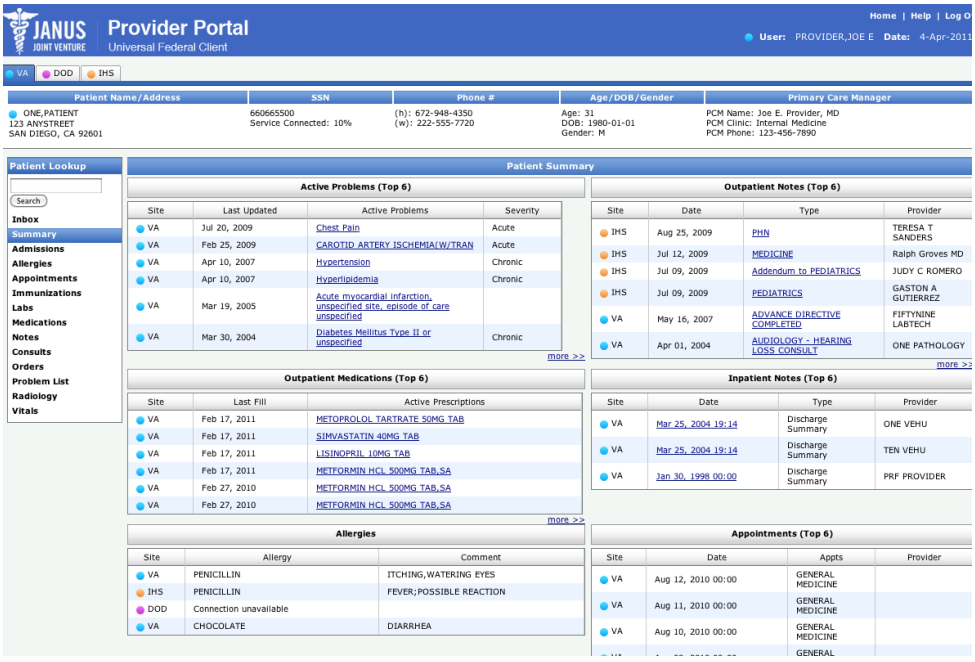

Figure 8: JANUS – The Universal Federal Provider Portal

The ability to provide role-based access control and flexible message filtering enabled the KMR team to repurpose the Inbox concept for the NHIN and VLER demonstrations.

Patient Medical Record

The second major application delivered by the team was the prototypical Patient Medical Record. It is based on the simple premise that a key to affecting long-term behavioral change is to engage our patients in a lifelong dialog educating them regarding not only there medical conditions, but also how to stay health and the challenges inherent in providing quality care. An informed consumer is a satisfied customer. Unfortunately, a personalized and collaborative healthcare experience is for many patients the exception rather than the rule. The Patient Medical Record (PMR) initiative enables patients to interact with their electronic health record and their providers in more dynamic and asynchronous ways.

The Patient Medical Record provides patients with a fully integrated collaboration environment having many of the same functional capabilities as an EMR. The tool provides basic, read-only functionality to access the entirety of a patient clinical record, similar in concept to a Personal Health Record implementations (PHR), although deliberately more complete in scope. It also delivers several more interactive workflow capabilities including medical device uploads, managing appointments, emailing providers, etc.

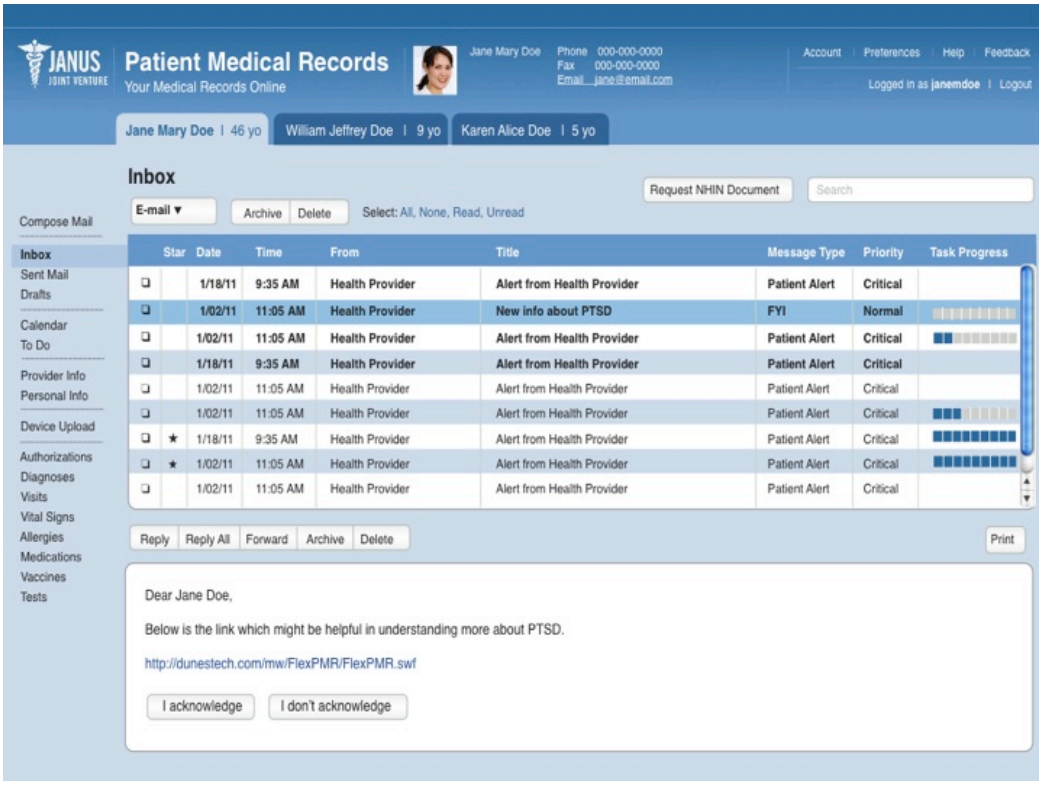

Figure 9: Patient Medical Record

The centerpiece of the Patient Medical Record application is a version of the Universal Inbox tailored to the needs of a patient. By providing them with the same capabilities that a medical professional enjoys, KMR was able to demonstrate how the tool could enable more dynamic, asynchronous interactions between patients and their primary care/subspecialty providers.

Android iAlert

Many of the functional scenarios developed recognized the need to deliver clinical decision support across different devices. These use cases highlighted that our workforce is mobile and cannot be assumed to have logged in a traditional desktop application. To ensure communications with this mobile workforce, KMR developed a CDS client for the Android cell phone platform. iAlert, as this prototype is called, refactored the Universal Inbox and tailored it for the provider or patient on the go. The app provides many of the same capabilities of the full-blown Universal Inbox including the ability to respond to clinical decision support messages requiring either message acknowledgment or other rule-defined actions. The application leverages the availability of the Android Telephony API to enable phone specific capabilities such as initiating telephone calls and e-mail/SMS text messages from within the app itself. The Android prototype provides a powerful demonstration for the numerous opportunities available to more immediately engage both patients and providers, and delivered valuable insight into how platform specific capabilities and form factor can be leveraged effectively.

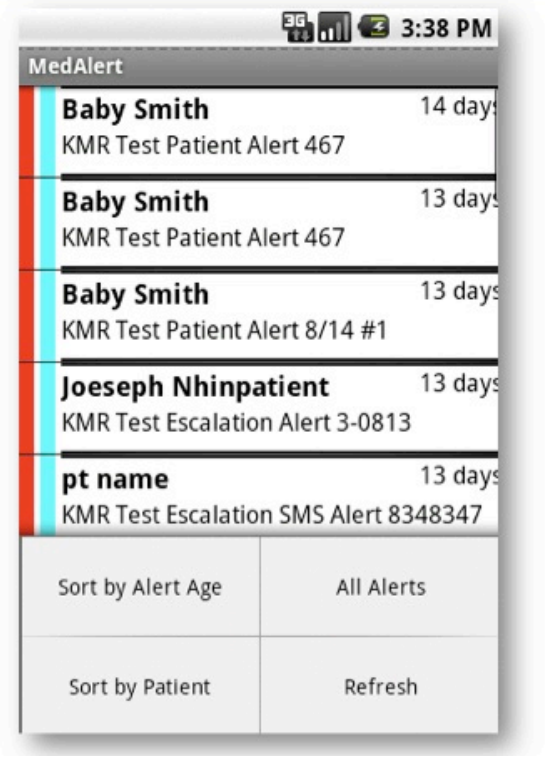

Figure 10: iAlert – Android Prototype for Mobile Devices

Deliverables are located in the "Project Source Code" folder on the Subject Data DVD.

10. Deliverable: System Administration Portal & Performance Monitoring Tool

The resource intensive KMR design necessitates a robust monitoring capability. Rule Monitoring Services is a second example of the value open access to source code enables. We were again able to extend the core capabilities of the stock commercial product with more extensive, standards-based, monitoring tools to provide the production manager of a large-scale CDS implementation to monitor the health and performance of large server farms with hundreds of rule sessions. These monitoring components are now embedded within the rules server and provide the ability to log almost any aspect of the production environment. Information regarding how many rules are deployed, the number of sessions, processor loads, resource availability, etc. can easily be logged for further analysis.

The Performance Dashboard delivers a flexible framework for exposing Rule Monitoring Service data need to manage large-scale CDS deployments. By providing visibility into data collected about runtime resource utilization, response times, number of new facts consumed, rules executed, recommendations generated, alerts ignored, actions taken, etc., an administrator will be able to review valuable operational information. Such data will prove invaluable in assessing the effectiveness of the CDS and ultimately justifying an organization's capital investment. The tool is intended to be equally valuable to the business process re-engineering team as it is to the operations manager.

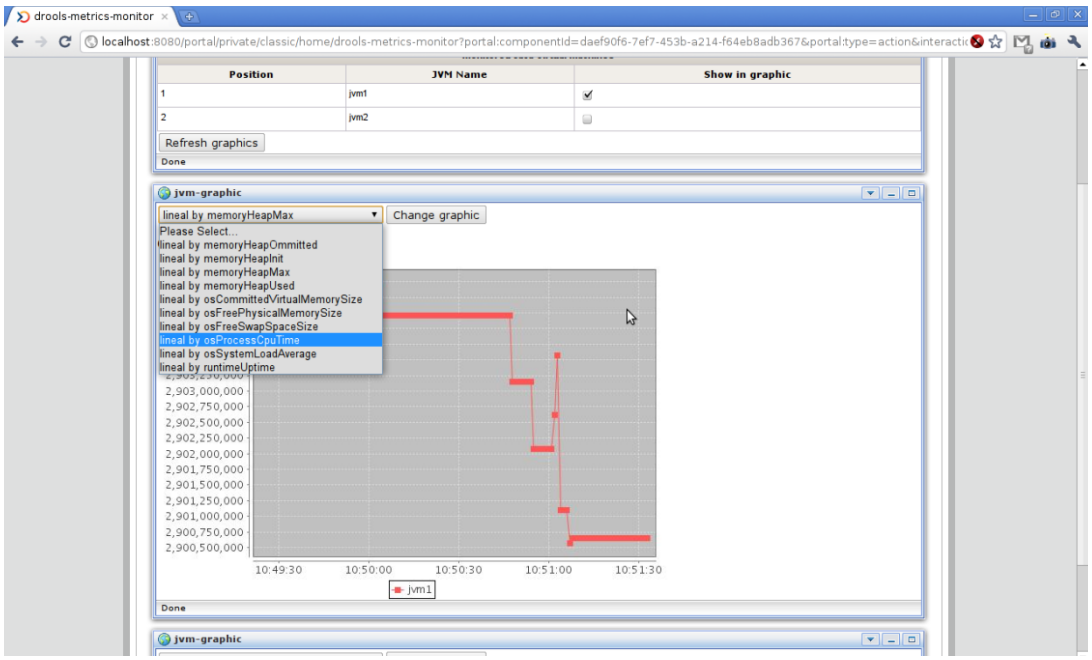

Figure 11: Run-Time Metric Portal

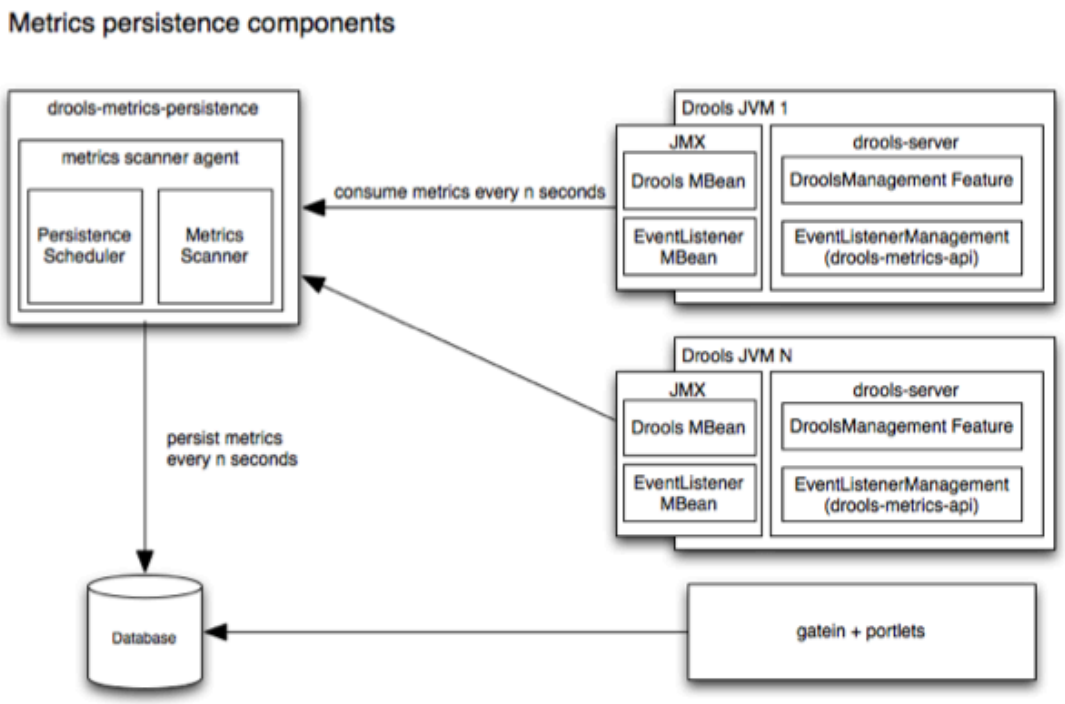

Figure 12: Rule Engine Monitoring Components

Deliverable is located in the "Project Source Code" folder on the Subject Data DVD.

11. Deliverable: Rules Workbench & Authoring Tools

The CDS Workbench provides the clinical domain expert with a graphical tool to author domain knowledge (rules, guidelines, other logic) using an intuitive application. The Workbench simplifies the creation of complex content by representing various clinical tests, activities, and procedures as "objects" in a drag-and-drop graphical editing environment. When the author is finished and wishes to deploy, the workbench compiles their work for storage and later retrieval by the run time system. Data elements exposed to authors for building rules are semantically constrained using the appropriate vocabularies stored in the KMR repository. When deployed within the runtime environment, the system can leverage this semantic meta-data to ensure that logical operations resolve with the highest degree of precision possible.

The Workbench has three panes. The left pane displays a menu of all the "facts" that the user is allowed to use within their rules. Each of these "facts" is a template reflecting clinical data, for example a lab or a medication, that is further defined using standard medical vocabularies exposed in the tool. The rule author sees a simple object representing a familiar clinical concept, but behind the scenes, KMR is ensuring that all components of the rule are perfectly aligned with the semantics and structural requirements of the facts that will be evaluated by the rule engine at run-time.

The center panel is the rule author's canvas. By dragging and dropping rule sub-components (facts, tasks, decision points, escalation events, etc.) onto this canvas, the

author can visually arrange them into a sequence that best reflects their innate understanding of the domain knowledge they are attempting to convey. The canvas is by its nature designed to reflect workflow considerations. However, the logic that controls the process flow of the guideline must also be incorporated. This is done using a pop-up editor that is exposed whenever the author double clicks on the graphical element representing a decision point in the guideline or rule. This pop-up enables the author to dive deep into the logic required by the guideline, but when complete, hide the editor and return to a more functional representation on the main canvas.

The right pane provides access to the meta-data and configuration parameters that a rule component needs to have defined. For example, if a user selects a Lab Object on the canvas, the right sided panels expose the appropriate LONIC search box so that the object and be further described as a CBC.

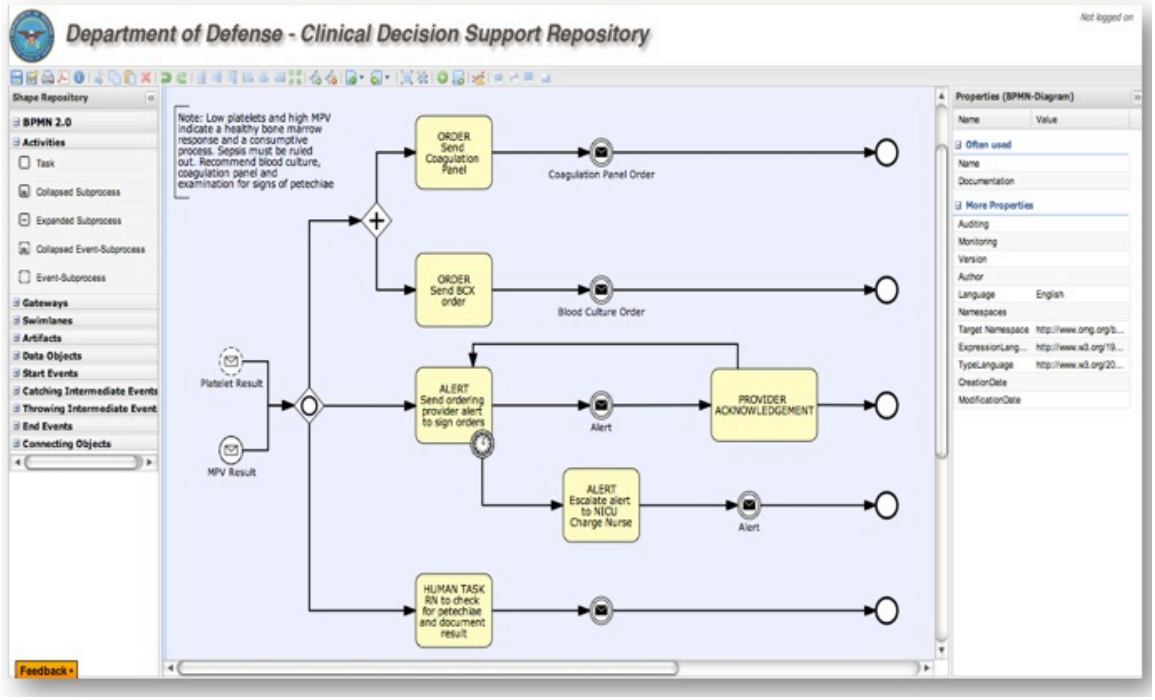

Figure 13: CDS Workbench

Deliverable is located in the "Project Source Code" folder on the Subject Data DVD.

12. Deliverable: Content & Executable Clinical Guidelines

The KMR services, tools, semantics, and data structures ensure that the full scope of rule logic and workflow descriptions can not only be accurately described, but also executed. We developed numerous use cases and implemented the actual rule logic to execute these clinical guideline recommendations within the context of real-time patient care.

Our initial work focused on the basic evaluation of laboratory data and illustrated that KMR had to be capable of performing two basic types of evaluation - independent and

dependent analyses. Independent evaluations do not need any information other than that contained within the triggering message and the domain knowledge embodied within the rule itself. The rule being executed may be very complex and even chained with other rules, but by definition no additional clinical data other than that contained within the message itself is required for execution. Dependent evaluations require additional information other than that contained within the triggering event message. Such evaluations require that KMR retrieve other clinical data from the EMR in order to execute. For example, one use case required that a provider be alerted when a platelet count was less than 80% of 6 previous platelet counts separated by a minimum of one week and all within 6 months of the current date. This statistical evaluation had to be done in real-time, with minimum impact on the performance of the transactional system, and with nearly instantaneous notification of one or more providers.

Later work focused on the processes by which a Rule generates a notification message to a recipient – these use cases helped define the standards-based interfaces (email, page, etc) and workflows that a notification might go through. For example, a three-step priority schema might use "normal", "low" and "high" where normal priority results are stored on the server for the ordering provider to retrieve as desired, low priority results are emailed, while high priority results warrant a page to both the provider and the clinic charge nurse. Rule workflows can also determine the time a notification may go unacknowledged before a message is escalated and to whom. Finally, alerts can be institutional or personal based on the role assigned to its author.

Our rule investigations culminated in a series of use cases designed to explore the ability to manage patient-specific rule execution in real-time and lead to our concept of a patient Cohort, the Cohort Service, and our patient-specific session architecture. A Patient Cohort is a group of *n* patients with similar attributes for which a given clinical evaluation applies. All clinical rules must be associated with at least one patient collect that scopes its execution. Patient Cohort can also be used symbolically to create logical conditions within a clinical rule. For example, one might define a "High-Risk Diabetic" collection as those patients with a HgA1c > 9. Once the basic collection is defined, it can be used to create a rule such as "if high-risk diabetic = true, then order HgA1c very six months, else every 12 months". Using patient Cohort allows rule authors to efficiently manage rule systems with thousands of individual rules as clinical definitions evolve or target values change.

Patients can be added to a cohort by applying other pre-configured rules that conveniently gather patients using attributes like PCM name, ward census or disease. If the patients are identified by a characteristic that is dynamic (e.g. location, disease, age) then they will be added and dropped from the collection automatically by the system according to whether they continue to satisfy the original criteria. For example, ward cohorts are composed of patients on a ward at execution time; clinic cohorts according to whether they have a current day or pending appointment.

If a Patient Cohort is defined by an attribute that is exclusively inpatient or outpatient, any rule applied to that collection will automatically be constrained to either inpatient or outpatient data respectively. For example, one might create a rule that enrolls patients in a renal disease registry if they have any two serum Creatinine values > 1.0. If that rule is applied to an outpatient cohort, patients with transient elevations of serum Creatinine

while hospitalized will not be added because their inpatient tests will be ignored for this particular rule.

A Patient Cohort has a scope determined by the role of its author. KMR supports as many roles as required, but is preconfigured with two basic ones to support institutional and personal degrees of granularity. It is import to understand that "institutional" always refers to a cohort that the organization, perhaps the health plan, has created and deployed. A general user cannot alter these. In contrast, "personal" collections were created by and are specific to a particular person –individuals have complete control over their own personal cohorts. Institutional cohorts are visible to all, while personal ones are private to the author.

Our use cases resulted in a system design able to accommodate all basic guideline requirements. We were not, however, able to develop specific executable guidelines for TBI, PTSD or Diabetes. The Geneva/NHRC team was focused on the engineering challenges to support the capabilities that such guidelines require. This was done successfully. The clinical implementation of these capabilities for TBI, PTSD or Diabetes guidelines was to be done with subject matter experts as part of our collaboration with the Indian Health Service and demonstrations slated for late summer 2010. Unfortunately, the IHS resources for this work were never made available and the demonstrations had to be dropped.

Without our clinical partner, the engineering team sought to use the National Quality Foundation eRecommendations for Diabetes as the source of clinical expertise for the diabetes guideline. We were indeed able to implement several of the best practice recommendations for diabetes, but the effort fell short of a comprehensive guideline demonstration as the eRecommendations released by NQF fell far short of the scope originally intended. Our search for implementable guidelines for TBI and PTSD also proved largely unsuccessful, not because the KMR infrastructure was found wanting, but because the semantics of the clinical data (psychometrics, cognitive function, quality of life, functional impairments, etc.) are in flux and largely unconstrained. Given that our entire infrastructure is driven by standardized terminologies and established data structures, the lack of such standards in several key areas of the TBI and PTSD domains ultimately proved to be a major impediment.

Creating new and proprietary specifications for TBI and PTSD was not within the scope of the KMR engineering team. Instead, we focused on demonstrating a capability central to the execution of any real-world clinical guideline…..the ability to reason over data stored across multiple organizations in real-time. We used a scenario involving a Federal employee (with data in DOD, IHS, and VA databases) who leaves the Federal enclave and receives care at a civilian facility during the later months of her pregnancy. During that time she develops gestational diabetes. Upon delivery she returns to the Federal enclave where the KMR system then evaluates the local Federal data, automatically requests and receives the distributed civilian data, and then reasons over this aggregate clinical information to determine what additional testing was needed. The KMR rule engine was able to make the diagnosis of true adult onset diabetes according to DoD/VA guidelines and then suggest additional care recommendations that took into account the aggregate information collated from across the NHIN – all duplicative tests and recommendations were thereby eliminated.

The source code for these executable guideline recommendations and rules is located in the "Project Source Code" folder on the Subject Data DVD.

13. Deliverable: Academic Presentation & Demonstration Of Runtime Deliverables

The Geneva Team provided several interim presentations and live demonstrations of those components of the KMR architecture as they became available. At AMIA 2009, we conducted a live demonstration of the core Decision Support Service and the first prototype of the real-time triggering of CDS rules and workflow. This was a daylong workshop with key representatives from Duke, UCLA, VA, IHS, DOD, AMIA and the American Medical Association.

In February 2010, we demonstrated the KMR Common Access Layer, Fact Extraction Service, and the Universal Inbox as full developed and DIACAP certified deliverables integrated within the AHLTA client and deployed as part of the San Diego VLER demonstrations. This Agency level demonstration received critical acclaim from both the clinical and technology communities, earning the team national recognition and awards. The VLER demonstrations proved unequivocally that the KMR design, semantics, metadata, and service oriented approach was not only scalable up to the national level, but was also applicable to a wide range of use cases, from document based health summary exchange to real-time clinical decision support for an individual patient.

At the March 2010 RSA demonstrations, in conjunction with our VA collaborators, we demonstrated how the KMR infrastructure could support rule-based encoding fine-grained access control down to the gene sequence level. In this demonstration, we used KMR semantics to define a series of patient directives (rules) describing when and how a patient's genetic information might be shared between organizations. In our demonstration use case, we illustrated how a patient might authorize the sharing of all their genetic results with the exception of any gene sequence known to be associated with schizophrenia. Additionally, we demonstrated how this supposedly "safe" information, once shared, could be further redacted to eliminate any information that at a later time became known to be a risk factor for schizophrenia. The use of our semantic infrastructure, the interoperability framework it helped created (FHA-CONNECT) and the KMR middle-tier services again demonstrated its robust and sophisticated design.

These interim demonstrations led up to the full KMR demonstration at HIMSS 2011 described below.

AMIA CDS Workshop November 2009	National Peer Reviewed Workshop
DOD-VA-Kaiser VLER February 2010	National Demonstration / Peer Review
RSA March 2010	National Demonstration / Peer Review
HIMSS March 2011	National Demonstration / Peer Review

14. Deliverable: Academic Review Of Current CDS 'State Of The Art'

In conjunction with the live demonstrations of KMR components, the Geneva Team participated in several academic panel discussions regarding the current "State-of the Art" in Clinical Decision Support. AMIA 2009 provided 2 peer-reviewed opportunities (a panel discussion and a workshop) to discuss the challenges facing CDS in general, and KMR in particular. CAPT Fry also had the opportunity to brief KMR requirements and design to all 20+ Federal groups that participate in Office of the National Coordinator's CDS Collaboratory initiative.

All these opportunities resulted in significant public and vendor exposure, and provided valuable affirmation that the system design was not only capable and scalable, but of considerable interest to the vendor community.

AMIA CDS Panel, November 2009	National Peer Reviewed Panel Presentation
CDS Collaboratory April 2009	Member, Office of the National Coordinator, CDS WG
CDS Collaboratory Update November 2009	Member, Office of the National Coordinator, CDS WG
CDS Collaboratory November 2009	Member, Office of the National Coordinator, CDS WG

15. Deliverable: Academic Panel On Future CDS Requirements

The Geneva Team participated in several academic venues regarding future requirements for Clinical Decision Support, at AMIA 2010 and at the International Medical Informatics Conference in Cape Town South Africa in Sept 2010. These were valuable opportunities to discuss the significant impediments facing our community as we strive for real-time rules and guidelines that can be shared between organizations.

The future state envisioned and discussed at these meetings highlight the need for appropriate standards and terminologies, several of which are currently poorly defined or lacking entirely. For example, the KMR team identified early on that the operational environment that a rule is designed to run in must be as well described semantically as the logic of the rule itself if a rule is to be shared safely between organizations. If a DoD pediatrician creates a diabetes management rule and makes that rule available to others, it would be inappropriate for an organization such a s the VA to take that rule and execute it within their environment and on their population of patients. In other words, the intended context of a rule is as important as the accuracy of the logic itself. His contribution is reflected in the *NQF CDS Report 2010* to which he contributed KMR experiences and insights as a Subject Matter Expert.

The KMR Team worked very closely with the VA and the OASIS TP20 Access and Control Security standards development team, because the challenges facing patient controlled authorization and disclosure are conceptually similar to the challenge of authorizing and executing clinical rules. The Team is particularly proud to have contributed so meaningfully to the *OASIS XSPA WS Trust Profile* standard, not only because it represents a significant academic achievement, but because it lies the

technological foundation for the future of Clinical Decision Support envisioned by the academic community and articulated in the KMR functional scenarios document.

A second major impediment to the widespread adoption of shareable, cross-organizational rules is the lack of a well-defined healthcare object model. HL7 has published a comprehensive Reference Information Model (RIM), but only has implementation specifications of sufficient definition to be used in messaging and CDA-based document exchanges. The lack of an unambiguous object implementation model prevents rule authors from creating rules that can be run in different organizations without needing a near-complete rewrite. We have the ability today to semantically define and share our rules - we do not have the ability to execute shared rules without costly refactoring.

To help address this significant impediment to the future envision for Clinical Decision Support, CAPT Fry participated in the HL7 Virtual Medical Record Working Group and the work that led to the publication of the *HL7 vMR Domain Analysis Model* that can be used to create a Virtual Medical Record over which KMR and similar systems can reason. He also co-authored the *Cross-Institutional CDS Data Needs Analysis* published by AMIA that year.

MEDINFO Presentation September 2010	International Peer Reviewed Presentation
AMIA Open Source Panel November 2010	National Peer Reviewed Panel Discussion
AMIA vMR Presentation November 2010	National Peer Reviewed Presentation
HL7 vMR Domain Analysis Model May 2010	Member, HL7 CDS Virtual Medical Record Working Group
NQF CDS Report 2010	Subject Matter Expert, National Quality Foundation
Cross-Institutional CDS Data Needs Analysis July 2010	Member, HL7 CDS Virtual Medical Record Working Group
XSPA WS Trust Profile April 2010	Member, OASIS - Healthcare Security and Access Control

16. Deliverable: Academic Review Of Completed KMR Project

In the summer of 2010, after the system requirements and design phases had been completed and the majority of the engineering validation had been completed, CAPT Fry met with the Principle Investigators leading three of the four largest academic CDS initiatives. He met with Dr. Richard Schiffman at Yale, Dr. Blackford Middleton at Partners Healthcare, and Dr. Ken Kawamoto at Duke, giving presentations on the full KMR scope and design to these pioneers and their respective university departments. In addition to these university based workshops, CAPT Fry presented KMR to the academic community at the International Medical Informatics Conference in Cape Town, South Africa in Sept 2010, and as a demonstrator invited by the Office of the National Coordinator at HIMSS 2011.

All these venues provided peer-reviewed opportunities to review the KMR project and to discuss the challenges facing CDS in general. The HMISS 2011 demonstrations validated the core concepts and architectural design of KMR. We successfully created,

deployed and executed clinical decision support rules across DOD, VA, and Indian Health Service electronic record systems. We demonstrated that using appropriate standards and carefully structured meta-data that the rules developed by one organization can be shared and executed in another. We illustrated that the same architecture used to successfully demonstrate NHIN and VLER interoperability can be used to satisfy countless other use cases, including real-time, distributed, and personalized clinical decision support.

The KMR project was received very well and there is considerable interest in several academic centers of excellence to fully evaluate the project in house once the code has been released to the public.

Yale, Harvard, Duke Workshops, July 2010	Peer Reviewed Workshop
MEDINFO Presentation September 2010	International Peer Reviewed Presentation
HIMSS Presentation March 2011	Office of the National Coordinator

17. Deliverable: Academic Outcome & Usability Evaluations

Our collaborative requirements process underscored the need for CDS to be exposed in a workflow sensitive manner that minimizes cognitive over load. This principle proved to be as important to patients as it is to providers. Both are typically juggling multiple tasks, commitments, priorities and deadlines. Under such circumstances, it is predictable that patients and providers become increasingly task focused and their ability to assimilate new information is limited. It is critically important that decision support recognizes the cognitive environment of the user and reduces unnecessary context shifting.

It became clear that meaningful and productive communication between patient and provider was essential in achieving measurable improvements in patient care, and that clinical decision support would prove to be a critical enabler. The requirement for a central workflow centric inbox to consolidate communications in an asynchronous manner began to emerge. This inbox was conceived of as a modular, plug and play component that can be embedded in a variety of different applications and settings. The inbox would not be restricted to merely e-mail or text based CDS alerts, but would be a type of multimedia “entertainment” center where a variety of message types could be viewed and responded to. Requirements for managing e-mail, CDS alerts, tasks, surveys, NHIN documents, and otherwise responding appropriately to different workflows began to emerge. Furthermore, this capability needed to be exposed within a variety of EMR applications, patient health records (PHR), and across a multitude of different devices including cell phones.

These requirements provided the team with cognitive framework for defining the usability concepts and workflows requirements that were envisioned during the development of our CDS scenarios. The team defined the need for patients to seamlessly integrate their personal calendars with their medical appointments, to access online schedules not only for their primary care providers, but also those of subspecialty consultants they had been referred to. The decision support infrastructure was envisioned

to safely and appropriately manage which available appointment slots were exposed and to whom. The desire to allow patients to request their own NHIN documents and to automatically include the returned data into the decision support process was also identified. The need to allow patients to access and view their own clinical data almost in real time was identified as both possible and appropriate if the required context and educational support could be provided.

All these workflows and capabilities were integrated into the KMR architecture and were intended to be fully implemented in our Indian Health Service demonstrations. Dr Doug Fridsma and the Arizona State University were to lead a Native American / University partnership to evaluate the usability and outcome improvements that well-integrated clinical decision support functionality would enable. Unfortunately, this pilot usability study was never done as the Indian Health Service was unable to free the resources it needed to security test and deploy the final KMR system. While deployment negotiations are still ongoing, the required pilot was not possible during the project's period of performance. As a consequence of this dependency, the KMR Project Manager/Liaison position was defunded, the ASU deliverables were re-scoped, the University's intended contributions to the KMR project severely curtailed and our usability concepts, workflows, and decision support innovations were never fully evaluated in a controlled study. This represents the only failure of the project, a failure for which the project plan, the deliverable schedule, and our emphasis on engineering innovation provided no adequate contingency plan.

Nevertheless, the Geneva Team did present KMR usability innovations and concepts at all the presentations, academic meetings, standards development meetings, and national demonstrations referenced above. We did the same at the community opportunities below and there is tremendous interest in leveraging our code and approach in many other initiatives as soon as it can be released through FHA-CONNECT under the Open Source BSD license or via FOIA.

UC Davis CDS Lecture December 2009	Peer Reviewed Workshop
TATRC PLR Presentation March 2010	Peer Reviewed Panel Presentation
ONC Presentation May 2010	Office of the National Coordinator
VA Terminology & Semantics June 2010	Peer Presentation
Consumer Choice Testimony July 2010	Congressional Hearing / Office of the National Coordinator
IHS Conference December 2010	Peer Presentation

KEY RESEARCH ACCOMPLISHMENTS

The KMR Project executed an ambitious and deliberately broad agenda heavily focused on engineering and the conceptual deliverables summarized above. Specifically, the following research milestones can be highlighted as most significant.

a) Functional Requirements

The project delivered a detailed use case and functional requirements document that was reviewed by subject matter experts during the 6-month analysis phase and presented

formally at the KMR CDS Workshop, AMIA 2009. This task was approached by conducting almost 6 months of weekly collaborative meetings with open-source community leaders, representatives from academic institutions such as Intermountain Health, Kaiser, Indian Health Service, VA and Arizona State University. The result was a comprehensive set of clinical scenarios and decision support business requirements representing perhaps the single most comprehensive collection of CDS use cases collected to date. These are also unique in that they articulate the particular requirements for cross-institutional rule and data exchange.

b) Canonical Model

The KMR canonical object model is a set of reference classes based on HL7 version 3 standards and vocabularies that normalize the structure and semantics of legacy data, thereby shielding the middle tier from the particulars of the enterprise data model. We believe this to be the most complete set of clinical object specifications available to the public as an HL7 RIM compliant, non-proprietary data model. HL7 3.0 RIM is both complex and difficult to understand - the KMR canonical model is a major stepping stone for developers attempting to implement a standards-based infrastructure. Until the medical community certifies one or more domain object models, the KMR object model is perhaps the best starting point for cross-organizational rule sharing and execution.

c) Individualized, Stateful CDS Session

The KMR approach ensures every patient is given their own instance of the rule engine into which individualized knowledge modules can be deployed. In addition, KMR leverages the concept of a stateful, in memory, persistent store of all historic facts that might be needed when evaluating new data. This design enables rules requiring statistical analyses, trend evaluations, or other complex operations to occur almost instantaneously. This patient specific, stateful approach is unique in the CDS literature, and opens the door to intelligent process and workflow management that is simply not possible when chaining multiple stateless rules together as is the current approach to automated clinical guideline development. As demonstrated at HIMSS 2011, the performance of the KMR system is outstanding and is believed to be scalable, given adequate memory resources, to millions of patient facts and knowledge bases with tens of thousands of rules.

d) Reference Run-Time System

The project delivered a live national demonstration of computable clinical guidelines being executed simultaneously against data from multiple clinical repositories, triggered using real-time HL7 messages, and incorporating Summary of Care documents received from the Nationwide Health Information Network. The demonstration used production Information Systems and EMR clients from VA, Indian Health Service and DoD, validating the applicability of the KMR Service Oriented Architecture for different Federal agencies. This reference system provides a comprehensive platform for refining and implementing CDS concepts and architectures. Being open-source, it sets the stage for multi-organization collaboration and research meta-analysis that is simply not possible when each team executes their research agenda on different systems, with different capabilities, semantics, and GUI metaphors. KMR essentially provides a

controlled reference platform for controlling confounding variables and incrementally delivering new functional capabilities.

e) CDS Architecture Contributions

Both the larger IT community and the Military Health Service have recognized the project's contributions to the national agenda. TATRC aligned its Advanced Technology Group Health IT portfolio around the KMR strategy and provided the additional funding and resources to take core architectural component to full production in the NHIN & VLER pilots. These demonstrations won TATRC national recognition as a winner of both Computerworld Laureate (2009) and CIO 100 (2010) awards. The NHRC/TATRC/KMR collaboration was recognized by MHS leadership for its contributions with numerous letters of commendation and appreciation.

REPORTABLE OUTCOMES

The project contributed significantly to the CDS literature though numerous national and international academic panel presentations, contributions to Standard Development Organizations, and peer reviewed publications. These are highlighted in the following section on reportable outcomes.

Presentations	Comments
CDS Collaboratory April 2009	Member, Office of the National Coordinator, CDS WG
CDS Collaboratory Update November 2009	Member, Office of the National Coordinator, CDS WG
CDS Collaboratory November 2009	Member, Office of the National Coordinator, CDS WG
AMIA CDS Workshop November 2009	National Peer Reviewed Workshop
AMIA CDS Panel November 2009	National Peer Reviewed Panel Presentation
UC Davis CDS Lecture December 2009	Peer Reviewed Workshop
TATRC PLR Presentation March 2010	Peer Reviewed Panel Presentation
ONC Presentation May 2010	Office of the National Coordinator
VA Terminology & Semantics June 2010	Peer Presentation
Consumer Choice Testimony July 2010	Congressional Hearing / Office of the National Coordinator
Yale, Harvard, Duke Workshops July 2010	Peer Reviewed Workshop
MEDINFO Presentation September 2010	International Peer Reviewed Presentation
MEDINFO Panel September 2010	International Peer Reviewed Panel Discussion
AMIA Open Source Panel November 2010	National Peer Reviewed Panel Discussion
AMIA vMR Presentation November 2010	National Peer Reviewed Presentation
IHS Conference December 2010	Peer Presentation
HIMSS Presentation March 2011	Office of the National Coordinator

Awards	Comments
Computer World Laureate Award 2009	Community Recognition Award
CIO 100 Award 2010	Community Recognition Award
Demonstrations	**Comments**
HIMSS April 2009	National Demonstration / Peer Review
RSA March 2010	National Demonstration / Peer Review
HIMSS March 2011	National Demonstration / Peer Review

Publications / Standards	Comments
HL7 vMR Domain Analysis Model May 2010	Member, HL7 CDS Virtual Medical Record Working Group
NQF CDS Report 2010	Subject Matter Expert, National Quality Foundation
Cross-Institutional CDS Data Needs Analysis July 2010	Member, HL7 CDS Virtual Medical Record Working Group
XSPA WS Trust Profile April 2010	Member, OASIS - Healthcare Security and Access Control
Morningside Collaboration August 2010	Peer Collaboration

CONCLUSION

The work accomplished during the KMR grant paints a clear picture of what is possible using current technology if the individual components can be orchestrated into a seamless whole. Its vision of a dynamic and workflow centric future where healthcare information technology enables and supports meaningful collaborations between patient, provider and machine is neither imposing nor de-humanizing. Instead, it illustrates that CDS technology, as a means to an end, is merely a tool in the services of society and the future it defines for itself. CDS will never replace provider or patient autonomy, nor will it alone be sufficient to deliver adequate, quality care. It can, however, contribute meaningfully to evidence-based, personalized care and relieve participants from some of the policy and administrative burdens they are increasingly subjected to.

The HMISS 2011 demonstrations validated the core concepts and architectural design of KMR. We successfully created, deployed and executed clinical decision support rules across DOD, VA, and Indian Health Service electronic record systems. We demonstrated that using appropriate standards and carefully structured meta-data that the rules developed by one organization can be shared and executed in another. We illustrated that the same architecture used to successfully demonstrate NHIN and VLER interoperability can be used to satisfy countless other use cases, including real-time, distributed, and personalized clinical decision support.

There is, however, much more to be done, especially in the areas of how workflow, process semantics, and alternative inference technologies can be orchestrated. For example, supply and demand forecasting in the MHS is predicated on the dual requirement of an accurate analysis of direct / purchased care resources within a community and equally accurate demand forecasting. This analysis must be sensitive to the limited flex capacity of rural and/or medically under-served areas. Unfortunately, current predictive models for simulating the effects of deployment-related conditions and projecting accurate resource requirements for the required follow-on care, especially regarding mental health disorders, are simply inadequate. It is not unusual for our treatment facilities to estimate demand upon return of a military unit using a fixed percentage of the number of troops deployed, regardless of where or how long they were in the field. The KMR infrastructure does not currently have the ability to call alternative

inference engines or technologies and, therefore, cannot manage the predictive models that might be required to seamlessly deliver such functionality to the end-user.

A second area where significant work remains to be done is research into the challenges that stateful orchestration of clinical guidelines implies. The KMR's workflow approach to CDS was extremely well received by functional domain experts; its process-centric metaphor fits well with the cognitive model that many providers use in daily clinical activities. Nevertheless, KMR does not currently take full advantage of standard workflow semantics and concepts. For example, it is common in workflow / process communities to distinguish between "delegating" and "transferring" a task. When delegating, a person assigns a task to another individual to complete - they are NOT reassigning responsibility which ultimately remains with them. When transferring a task, the user IS reassigning ultimate responsibility, for example, when they turn over care of an inpatient to the call team at the end of the day. This distinction becomes increasingly more important when workflows are linked together, either in series or in parallel, through time. There is no automated clinical practice guideline in existence today that even attempts to encode and manage task responsibility. KMR provides an architecture that for the first time can realistically begin to investigate how to both articulate and to encode such concepts.

A final area where more investigation is needed is in the terminology services required to truly share CDS artifacts. Rule interoperability in KMR is achieved through the rigorous application of HL7 standards and semantics. It is, however, an architecture driven largely by convention. When a rule written to utilize one vocabulary (e.g. SNOMED) is presented with "facts" that utilize a different vocabulary (e.g. ICD9), it will refuse to execute until either the rule or the data is semantically transformed. This possess some very interesting questions regarding when it is appropriate to change the terminology of the data, and when the rule itself should be refactored to accommodate the desired semantics. Again, KMR provides an ideal platform for investigating these difficult questions.

Acronyms

Acronym	Definition
aCPG	Automated Clinical Practice Guideline
AMIA	American Medical Informatics Association
API	Application Programming Interface
CAL	Common Access Layer
CDR	Clinical Data Repository
CPG	Clinical Practice Guideline
DIACAP	DoD Information Assurance Certification & Accreditation Process
DoD	Department of Defense
DSS	Decision Support Service
DT&E	Development, Test, and Evaluation
EMR	Electronic Medical Record
FHA	Federal Health Architecture
HIEM	Health Information Exchange Message
HIPAA	Health Insurance Portability and Accountability Act of 1996
IHS	Indian Health System
HIT	Health Information Technology
HL7	Health Level Seven
HTTP	Hypertext Transfer Protocol
ID	Identification
IHE	Integrating the Healthcare Environment
JSON	Java Script Object Notation
KMR	Knowledge Management Repository
MHS	Military Health System
MRMC	Medical Research and Materiel Command
MTF	Military Treatment Facility
NHIN	Nationwide Health Information Network
NQF	National Quality Foundation
OASIS	Organization for the Advancement of Structured Information Standards
PCM	Primary Care Manager
PDF	Portable Document Format
PHI	Personal Health Information
PHR	Personal Health Record
PMR	Patient Medical Record
POC	Point of Contact
PTSD	Post Traumatic Stress Syndrome
QA	Quality Assurance
RIM	Reference Information Model
SOA	Service Oriented Architecture
SOW	Statement of Work
SQL	Structured Query Language
SRS	Software Requirements Specification
TATRC	Telemedicine and Advanced Technology Research Center
TBI	Traumatic Brain Injury
UDDI	Universal Description, Discovery and Integration
UI	Universal Inbox
UI	Universal Inbox
VA	Veteran Administration
VLER	Virtual Lifetime Electronic Record
vMR	Virtual Medical Record

Key Personnel

Name	Comments
CAPT Emory Fry	Naval Health Research Center, San Diego
Kim Pham	TATRC
Jerry Goodnough	TATRC
Tuyet Nguyen	Geneva Foundation
Chrisjan Matser	Geneva Foundation
Duane Decouteau	Geneva Foundation
Mike Wright	Geneva Foundation
Fateh Hilal	SOADex, Inc

Requirements Working Group

Name	Organization
Theresa Cullen	Indian Health Service
Stan Huff	Intermountain Health
Craig.W.Robbins	Kaiser Permanente
Robert Greenes	Arizona State University
Steve Steffensen	TATRC
Clayton Curtis	Veteran Affairs
Doug Fridsma	Arizona State University
Mary Goldstein	Veteran Affairs
Nathan Hulse	Intermountain Health
Andy Wiesenthal	Kaiser Permanente
Peter Haug	Intermountain Health
Goodnough Jerry	TATRC
Pham Kim	TATRC
Duane DeCouteau	Geneva Foundation
Matser Chrisjan	Geneva Foundation
Jeff Kriseman	Arizona State University
Howard Hayes	Indian Health Service
Hemant Shah	Henry Ford Helth System
Ken Kawamotto	Duke University
Shabeer Sayid	SOADex
Hilal Fateh	SOADex